

# Cisco IOS QoS

## Frequently Asked Questions

A blue square containing a white stylized globe graphic with latitude and longitude lines.

### NBAR

Q. What is NBAR?

A. Network-Based Application Recognition is a classification tool that can identify traffic up to the application layer. Its protocol discovery feature provides per-interface, per-protocol, and bidirectional statistics for each traffic flow transiting an interface.

Q. When will NBAR be supported on the Cisco 7500?

A. As of now, NBAR is supported on the Cisco 7100, 7200, and 7500 (Route Switch Processor [RSP]-based) platforms. Contact Scott Frisby for progress on Versatile Interface Processor (VIP)-based dNBAR support.

Q. What are PDLMs?

A. Packet Description Language Modules make classification possible. PDLM is a list of all the protocols that can be recognized by NBAR.

Q. Are PDLMs customizable for new applications?

A. For now, PDLMs are not customizable. New PDLMs for emerging applications are released by Cisco Systems. Please contact Scott Frisby to get support integrated for new applications.

Q. Where can I find Napster PDLM?

A. <http://www.cisco.com/cgi-bin/tablebuild.pl/pdlm>

Q. How many match statements does NBAR support per protocol?

A. NBAR supports up to 24 concurrent URLs, HOSTs, or Multipurpose Internet Mail Extensions (MIME)-type matches. It means that in your configuration you cannot have more than 24 statements that match on URL or MIME, regardless of whether they are in one class-map or spread across multiple class-maps.

Q. What is the expected memory utilization per-flow by NBAR during the inspection process?

A. NBAR uses about 150 bytes of DRAM for each flow that requires stateful inspection. 1 MB DRAM will support up to 5000 concurrent flows. If more memory is needed, it expands its memory usage in increments of 200 to 400 kb.

Q. Is NBAR supported within Modular QoS CLI?

A. Yes. NBAR for classification can be invoked using the **match protocol** option under class maps in modular QoS CLI (MQC). The protocol discovery needs to be enabled using **ip nbar protocol-discovery** under the individual interface of interest.

Q. NBAR supports subport classification; what does this mean?

A. It means that NBAR can classify application traffic by looking beyond TCP/ ports. It looks into the TCP/UDP payload and classifies traffic based on content within the payload.

Q. Are H.323 and SIP supported by NBAR?

A. Yes.

Q. Where can I go for more information?

A. Check the following URLs:

[http://www.cisco.com/warp/public/cc/so/neso/ienesv/cxne/nbar\\_qp.htm](http://www.cisco.com/warp/public/cc/so/neso/ienesv/cxne/nbar_qp.htm)

[http://www.cisco.com/warp/public/779/largeent/issues/app\\_net.html](http://www.cisco.com/warp/public/779/largeent/issues/app_net.html)

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121newft/121t/121t5/dtnbar.htm>

## Marking

Q. When do packets get marked? Before or after queuing?

A. Packets get marked before queuing.

Q. Is Cisco Express Forwarding (CEF) required for marking/remarking packets?

A. CEF is required if you are using the **SET** option within MQC to mark packets. If packets are marked using CAR or the Class-Based Policer, CEF is not required.

Q. Can locally generated packets on a router be marked?

A. Locally generated packets or packets destined to the router never get CEF switched, so they cannot be marked via the **SET** option. Packets can be marked via the outbound Class-Based Policer or CAR, if this feature is enabled. Remember that dCEF is required for all of the QoS features on the Cisco 7500 platform if running in distributed mode.

Q. Can packets be remarked in priority class (in LLQ)?

A. Yes. Packets can be remarked in any of the classes using the **SET** command. You can also use policing to remark packets within a priority class, but this is not desirable because the Policer enforces the maximum rate.

Q. Which field of a packet is used for marking or remarking packets? What values are available? What configuration method is used to accomplish this?

A. It depends on if you are marking packets at Layer 2 or Layer 3. For layer Layer 2 class of service (CoS), 8 levels of priorities can be set using the 3-bit TAG field in the 802.1Q/P frame or 3-bit Inter-Switch Link (ISL) header in the ISL frame. For IP packets, 8 levels of priorities can be set using the 3 bits of the type-of-service (ToS) field if you are marking with IP Precedence. If you are using Differentiated-Service-Code-Point (DSCP) values, 64 levels of priorities can be set using 8 bits from the ToS field. Eight priority levels for Multiprotocol Label Switching (MPLS) are set using 3 experimental bits. For ATM, the Cell-Loss-Priority (CLP) bit is changed from 0 to 1. All of these can be configured using QoS Policy Manager (QPM) or MQC. To alter the Frame Relay Discard-Eligible (DE) bit, you will need to utilize interface-level EXEC commands. Support for this via MQC is planned in 12.2T.

Note: Please refer to the Diffserv FAQ for additional information.



## MQC

Q. What is MQC?

A. MQC stands for Modular QoS CLI. It's a new model that supports network-wide policy-based QoS configuration and deployment. Diffserv traffic conditioner is composed by using MQC. MQC consists of three components: class-maps, policy-maps, and service-policy. Classification is defined under class-maps, policies are constructed for each class or combination of classes under policy-maps, and policies are attached to a desired Diffserv node by using the **service-policy** option.

Q. What are the differences between MQC and Class-Based Weighted Fair Queuing (CBWFQ)?

A. CBWFQ is a queuing policy that can be configured via MQC. Introduced in 12.0.5T, CBWFQ was the only queuing method that was introduced in MQC. Additional components of Diffserv traffic conditioner such as shaping, congestion avoidance, and LLQ are supported in the later Cisco IOS<sup>®</sup> Software versions.

Q. How many classes can be configured on a router?

A. 256.

Q. How many classes are supported under one policy?

A. Since you can define a maximum of 256 classes, you can define up to 256 classes within each policy if the same classes are reused for different policies. If different sets of classes are used for different policies, then the maximum number of classes used among all the policies cannot exceed 256. In this case, a total of 256 classes are supported per policy or per multiple policies if the policy does not include CBWFQ. In other words, if you have two policies, the total number of classes from both policies should not exceed 256. If a policy includes CBWFQ (in other words, if a policy contains a “bandwidth” [or “priority”] statement within any of the classes), the total number of classes supported will reduce to 64. If you have two policies, one containing CBWFQ and the other without CBWFQ, the total number of classes between two policies should not exceed 256, given that a policy with CBWFQ can have up to 64 classes and the other policy can have up to  $256 - 64 = 192$  classes.

Q. Is there a limitation on the number of interfaces a policy can be applied to?

A. No. A policy without CBWFQ can have up to 256 classes and can be applied on multiple interfaces.

Q. How many policies can I apply per interface?

A. Two—one in the incoming and the other in the outgoing direction. This can be the same policy or different policies. Make sure the QoS allocated is consistent and does not overlap, as well as policy-containing queuing/shaping is not applied in the incoming direction.

Q. Can the same class-maps be used with different policies?

A. Yes.

Q. I have configured the policy but it is not activated on an interface. What should I do?

A. Make sure the policy is attached to an interface of interest. Just creating a policy is not enough.

Q. Can I use the same class-maps within different policies?

A. Yes.

Q. Is MQC supported only in the CEF switching path?

A. MQC is supported in the CEF, fast, and process switching paths, depending on the platforms. For example, CBWFQ is supported in the fast and process switching paths on the Cisco 7200.

Q. When I configure the service-policy output statement under the ATM interfaces (either main or subinterfaces), the router takes the command but does not show it in the running configurations. The same problem persists on the Frame Relay subinterfaces. Are these interfaces supported, or do we need to configure them on a per-VC basis?

A. Service-policy needs to be configured on a per-VC basis. For the ATM case, the service policy needs to be attached to the ATM PVC. For the Frame Relay case, you need to have 121(2) T or higher and attach the service-policy to the map-class and then attach the map-class to the DLCI/subinterface.

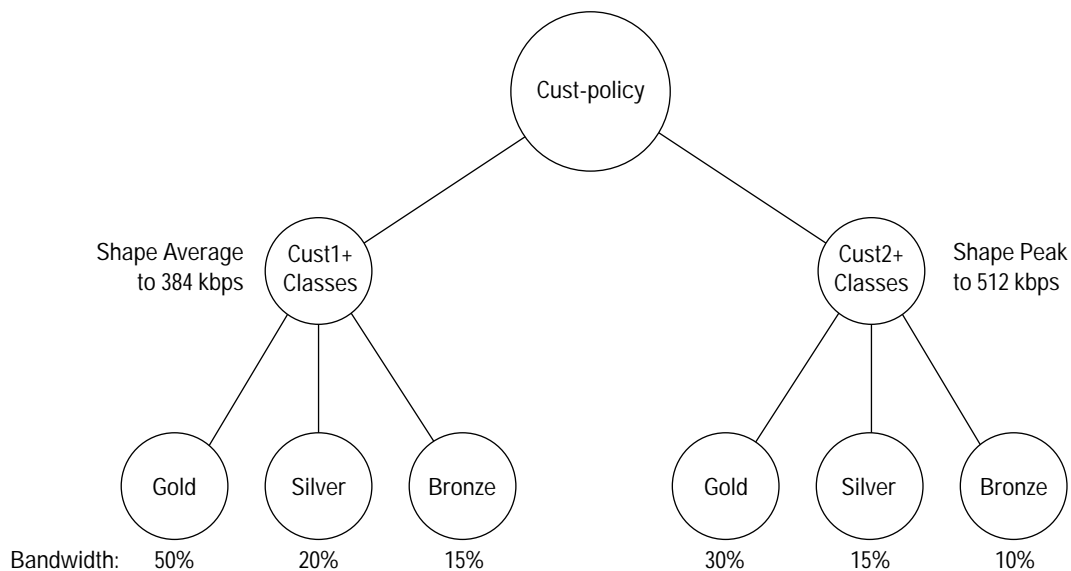
Q. If packets are marked or remarked within a policy using the **set** option, will the packet get treated based on the original DSCP value or on the new DSCP value?

A. This is platform dependant. On Cisco 7200 and lower-end platforms, policies are applied based on the old DSCP value. On the Cisco 7500, policies are applied based on the newly marked value.

Q. What is hierarchical policy?

A. The hierarchical policy mechanism lets you define policies within a policy. In order to configure this, define child policies, define a parent policy, and then apply the child policy to a parent policy by using the **service-policy <child policy name>** command under a class.

**Figure 1** Hierarchical Policy



As shown in Figure 1, *cust1-classes* and *cust2-classes* are child policies, whereas *cust-policy* is a parent policy. Any QoS you define within a child policy applies only to the traffic within that child policy. Any QoS you define within a parent policy applies to all the children policies on an aggregate basis. This sets the aggregated upper limit for the QoS resources that can be utilized by all the children policies.

Any-to-any QoS feature combinations are supported on the Cisco 7500 platform. For the Cisco 7200 and lower, the any-to-any QoS combination is not supported between parent and child policies. For example, shaping has to be at a parent level to support queuing at the child level. Hierarchical policing is supported. Hierarchical shaping is not supported yet (...12.1T. Check 12.2T for update). Additional supported combinations are given in Table 1:



**Table 1** QoS-Supported Feature Combinations

QoS @ parent level >	Bandwidth & Shaping	Bandwidth	Shaping	Police
QoS @ child level >	Bandwidth	Police	Police Priority Bandwidth	Police

Q. Where can I go for more information?

A. Check the following URLs:

<http://www.in-istm.cisco.com/qos/mqc.html>

WFQ/CBWFQ/LLQ

Q. Is it possible to activate CBWFQ for two serial interfaces bundled using virtual templates?

A. Yes. Apply the policy to the virtual template.

Q. Why is the transmit delay of prioritized traffic much worse with dWFQ (on the Cisco 7500-PA-MC-4T1) than CBWFQ on the Cisco 7200?

A. The LLQ feature is required to support low-latency behavior on a *priority* class, particularly on low-speed links. Distributed LLQ is available in Versions 12.1E and 12.0S, but not in Version 12.1.5T. Standard dWFQ will not provide low-latency behavior at low speeds (though at higher speeds [for example, 512 Kbps and above] it is typically acceptable).

Q. Are Priority Queuing (PQ) or Custom Queuing (CQ) supported at the ATM subinterface level?

A. Neither PQ nor CQ are supported on any ATM interface. It is recommended to use CBWFQ or LLQ per VC (they are not supported for subinterfaces, however).

Q. Is CBWFQ supported on BVI interfaces?

A. BVI is a logical interface, so it is not possible to do any queuing on it; queuing must be done on physical interfaces. If there is a need for this feature, please create a **sys-wish**.

Q. Is CBWFQ supported on ATM SVCs?

A. Per-VC CBWFQ is supported for PVCs only.

Q. Is CBWFQ/LLQ supported on virtual template interfaces?

A. There is no active queuing or any other functionality on virtual template interfaces. It holds the configuration only for virtual access interfaces. So, yes, policy should be applied to a virtual template, and associated QoS will be applied to virtual access interfaces.

Q. With WFQ there is a default maximum of 256 conversations. What happens if we exceed the maximum?

A. The flows just “collide” and the same entry is used in the hashing table. In this situation, several flows will share a subqueue. The number of subqueues cannot be increased indefinitely.

Q. The default number of WFQ dynamic queues is changed in Cisco IOS 12.1, so that the default value is automatically adjusted according to the bandwidth. The former default value was 256. What is the reason for this modification?

A. It was done as part of CSCdm26683, mainly to conserve memory on platforms that have numerous low-speed interfaces that do not benefit from having a large number of queues.

Q. What is the equation used to calculate the number of queues based on bandwidth per interface? Per VC? Is it linearly proportional?

A. It's 256 queues for an interface with 2M or higher bandwidth.

**Table 2** QoS Feature Combinations**For Interfaces:**

bandwidth<=	64k	128k	256k	512k.....	2M or more
queues-	16	32	64	128	256

**Table 3** VCS**For VCS:**

bandwidth<=	64	512	2000	8000	>than 8M
queues -	16	32	64	128	256

Because there can be a larger number of VCs than interfaces, it is memory-wise effective to have a fewer number of queues per VC than per interface.

Q. How is routing update traffic treated when using CBWFQ? Do we need to assign bandwidth for routing updates, or is it treated with highest priority automatically?

A. Open Shortest Path First (OSPF), Routing Information Protocol (RIP), Enhanced Interior Gateway Routing Protocol (EIGRP) (non-TCP) hellos, and keep-alives are marked crucial and get queued in the `pak_priority` queue. By default, routing updates, BGP keep-alives (TCP based) get marked with IP Precedence value 6 (higher priority) gets queued in the class `class-default` along with other unidentified traffic. Flow-based WFQ operates on all the traffic present in this class. Up to 25 percent of the interface bandwidth is available to this class. If more bandwidth needs to be guaranteed for the class `class-default`, then allocate it using the **bandwidth** command.

Q. Do we support per DLCI CBWFQ on a VIP-based system on an MPLS Provider Edge (PE) (facing the customer)?

1. Which version of IOS Software?

2. Is this implemented with a Frame Relay map class as with smaller systems or with a service-policy?

A. Yes, dCBWFQ, dLLQ, DTS, and FRF.12 are all supported on the VIP on a per-DLCI basis.

1. The Cisco IOS versions are 12.1(2) E1 or higher, and the upcoming 12.1.5T.

2. Yes, a service-policy is attached to the map-class. The only difference is that DTS CLI will replace the Frame Relay traffic shaping (FRTS) CLI.

Q. Is CBWFQ supported per VC with FRTS?

A. Yes, you may apply CBWFQ on top of an FRTS policy on a PVC. The feature will be available in 12.1(5) T.

Q. What switching path does LLQ run in?

A. LLQ runs in the CEF path; dWFQ will run in the dCEF path.

Q. Are dWFQ or dCBWFQ supported per Frame Relay VC?

A. Yes. CBWFQ may be applied on top of an FRTS policy on a PVC in 12.1.5T and later.

Q. Can CBWFQ be used with flow-based WRED on an interface? The documents show only CBWFQ + WRED.

A. No, CBWFQ cannot be used with flow-based WRED.



Q. Does CBWFQ support non-IP traffic? Can a non-IP access list be used with **match access-group** to classify Internetwork Packet Exchange (IPX), for example, into a class?

A. Yes. CBWFQ supports all protocols, and IPX access-lists can be used for classification via the match option.

Q. Does WFQ (no VIPs, no class-based) use IP Precedence bits in the queuing decision?

A. In WFQ, IP packets are classified into flows using the following criteria:

- ToS bits in the IP header
- IP protocol type
- Source IP address
- Source TCP or UDP socket
- Destination IP address
- Destination TCP or UDP socket

The combination of weight and the packet length is used to determine the departure time of the packet. In general, the higher the precedence, the lower the weight and the higher the bandwidth share.

However, IP Precedence is not used in deciding the hash queue number for the packet. The hash queue is decided based on the 5-tuple (src-ip-addr, dest-ip-addr, src-port, dest-port, protocol).

As packets are sorted into the queues, they are given weights that are, in turn, used to calculate a sequence number. The sequence number determines the order in which packets are dequeued and transmitted. These weights are calculated using the following formula:  $\text{Weight} = 4096 / (\text{IP Precedence} + 1)$

So, yes, IP Precedence does influence the decision as to when the packet will be scheduled.

Q. Does dWFQ (VIPs, but no class-based) use IP Precedence bits in the queuing decision?

A. The VIP scheduling algorithm does not compute sequence numbers/weights as in the Cisco IOS Software implementation; ToS-based dWFQ will account for IP Precedence and put it in the appropriate queue. The queues are serviced based on the weights assigned to them. So with ToS-based dWFQ, IPP determines which queue the packets get into, and then it is up to you to configure the share of bandwidth you want to give for traffic in the queue.

Q. In CBWFQ, the default class (class-default) gets the remaining bandwidth based on its weight after the reserved bandwidth gets distributed to other classes. Are the weights based on bandwidth specified for that class? If so, what happens if the default class is specified without a bandwidth? Shouldn't its weight be zero?

A. In default class, the weights are assigned to the packet based on their precedence in flow-based WFQ. These weights are far more than that of the bursty class weight, so the bursty class will get a large amount of unused bandwidth but the default class will not get starved.

Q. With CBWFQ, do all classes get WFQ treatment?

A. It is first-in, first-out (FIFO) within the class unless you define WRED. The only class that gets flow-based WFQ treatment is the class *class-default*. So between class platinum, gold, and bronze, WFQ is applied, but within class platinum or gold, FIFO is applied.

Q. In per-VC CBWFQ over ATM, given two classes, voice and data, data traffic fills up the queue until the congestion occurs. When congestion occurs and CBWFQ gets invoked, there might be several numbers of data packets present in the FIFO queue already, a scenario that could delay the transmission of voice packets. How can you fine-tune to guarantee minimum latency to the voice packets?

A. There is a default **tx-ring-count** of 40 associated with each ATM VC that indicates the number of outstanding packets a VC can have on the segmentation and reassembly (SAR) chip. Reduce the **tx-ring-count** number to reduce the latency for voice packets. (Should the hold queue be tuned for the same concern on other interfaces? No, the hold queue should not be modified, because it will not have the same effect).

Q. What is the maximum aggregated bandwidth that you can specify for CBWFQ classes, including the PQ class? Why? How can you guarantee bandwidth to class default?

A. By default, 75 percent of the actual bandwidth on an interface; the remaining 25 percent is available for the default class, but all of the 25 percent is not guaranteed to the default class for the Cisco 7200 and lower-end platforms. You will need to guarantee minimum bandwidth by configuring “bandwidth” explicitly if you need to assure minimum bandwidth for the default class. If there is no traffic in the default class, then other classes should be able to use this bandwidth accordingly. The default of 75 percent can be altered using the **max-reserved-bandwidth** command.

The router will give a warning message if you try to overprovision. For example, the following message was received when 75 percent of the bandwidth was reserved for other existing classes and the configuration tried to reserve more for the PQ class:

“If Serial4/1 class platinum requested bandwidth 200 (kbps) Available only 0 (kbps)”

Be aware when you use the **percent** option to configure the bandwidth. No warning message is sent if you overprovision. If you configure 10 percent and all the bandwidth is allocated to other classes, 10 percent of the 0 available bandwidth will be allocated!

```
7200-UUT#sh policy-map interf s4/1
Serial4/1
```

Service-policy output: TEST (2264)

```
Class-map: platinum (match-all) (2265/6)
  0 packets, 0 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
  Match: ip dscp 46 (2267)
  Weighted Fair Queuing
  Strict Priority
  Output Queue: Conversation 264
  Bandwidth 1511 (kbps)
  (pkts matched/bytes matched) 0/0
  (pkts discards/bytes discards) 0/0
```

```
Class-map: gold (match-all) (2269/2)
  0 packets, 0 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
  Match: ip dscp 10 12 14 (2271)
```



## Weighted Fair Queuing

Output Queue: Conversation 265

Bandwidth 10 (%) Max Threshold 64 (packets)

(pkts matched/bytes matched) 0/0

(pkts discards/bytes discards/tail drops) 0/0/0

Class-map: class-default (match-any) (2273/0)

0 packets, 0 bytes

5 minute offered rate 0 bps, drop rate 0 bps

Match: any (2275)

0 packets, 0 bytes

5-minute rate 0 bps

Q. What is the maximum bandwidth that can be reserved for the PQ within CBWFQ?

A. By default, 75 percent. If you allocate 75 percent of the bandwidth only for the PQ, then no minimum bandwidth can be guaranteed to other classes. The rest of the classes, including the default class, will end up sharing 25 percent of the bandwidth. If you need to alter the default values, use the **max-reserved-bandwidth** option.

Q. When does the default class (class-default) get created?

A. The class default preexists, although it may not be apparent in the **show policy-int** output log until a policy is applied to an interface.

Q. Is bandwidth policed per class or per queue?

A. Bandwidth is policed per class.

Q. Is bandwidth policed per flow (within a class) or per class?

A. Bandwidth is policed per class and not per flow.

Q. Does any of the fancy queuing get invoked in the absence of congestion?

A. No congestion, no fancy queuing; only FIFO will be used.

Q. Is CEF required for queuing (WFQ, CBWFQ, LLQ) to operate on locally generated traffic?

A. No.

Q. Can packets be remarked in priority class (in LLQ)?

A. Yes. Packets can be remarked in any of the classes by using the **set** command. You can also use policing to remark packets within a priority class, but this is not desirable because the Policer enforces the maximum rate.

Q. What is a significance of the burst attribute in the priority queue of LLQ?

A. This allows you tune Tc (burst interval). The default is 200 ms. Remember, Tc = Committed Burst (Bc)/committed information rate (CIR). Remember that Bc is in bytes (32-2000000), and CIR is in kbps (8-2000000). If your application requires less than 200-ms transmit time, then set Bc accordingly. Notice that there is a limit for the minimum required Tc. This varies, based on CIR, and is half the size of CIR (exclude kbps in the measurement). For example, if your CIR is 2000 kbps, then the minimum Bc size you can set is 1000. If your CIR is 1000 kbps, then the minimum Bc you can set is 500.

Q. How many priority queues are supported? Can more than one class do priority queuing? In case of multiple priority queuing classes, which class is Tc derived from?

A. There is only one priority queue. Even if you define priority queue within multiple classes, all the traffic from different classes will get queued into one queue. The advantage is that each class can be policed at a different CIR. Each policer will have its own token bucket. Tc will depend on specified CIR and the Bc value per policer. In other words, packets from one class could be inserted into priority queue at a different interval than the packet from the second priority class.

Q. Why are **random-detect**, **bandwidth**, and **queue-limit** not supported in the priority class?

A. WRED is intended for TCP traffic that is not delay sensitive; in addition, if dropped, the information will be retransmitted. Priority queue within LLQ is intended for delay-/jitter-sensitive traffic that loses value if delayed, and information is lost if the packets are dropped because the source is very unlikely to retransmit the same information. Allocating bandwidth by using the **bandwidth** command assures minimum guaranteed bandwidth. In LLQ, this is built into the **priority** command. **Queue-limit** is associated with WFQ, which is not the queuing technique used in the strict priority queue (class).

Q. Can Multilink PPP (MLP) run in conjunction with LLQ?

A. Yes, since Cisco IOS Version 12.0.7T. From the Point-to-Point Protocol (PPP) perspective (including MLP), you are either using WFQ or you aren't. The internals of how WFQ chooses to dequeue packets are entirely transparent to MLP. In case of dialup connections, FIFO is activated and fancy queuing is disabled as soon as an interface becomes part of a multilink bundle. In 12.2T, fancy queuing will be reenabled once again when the interface is removed from the multilink bundle.

Q. Are there any QoS processes that give strict priority to voice over IP-over-Frame Relay (VoIPoFR) traffic for one subinterface (int s 0.2) over data traffic on another subinterface (int s 0.1)?

A. Yes, there is a feature called Frame Relay PVC Interface Priority (PVCPIPQ) in 12.1.1T. Frame Relay PIPQ allows you to configure a PVC transporting voice traffic to have absolute priority over a PVC transporting signaling traffic, and a PVC transporting signaling traffic to have absolute priority over a PVC transporting data. There are three steps to implementing it:

1. Configure PVC priority in a map class.
2. Enable Frame Relay PIPQ and set queue limits.
3. Assign a map class to a PVC.

For example:

```
interface serial0.1
 encapsulation frame-relay
 frame-relay interface-queue priority 10 20 30 40
 frame-relay interface-dlci 101
 class high_pvc
 ...more DLCIs...
!
! Give PVC priority in the map-class
map-class frame-relay high_pvc
 frame-relay interface-queue priority high
```

Q. Is large-scale PPP over ATM (PPPoA) supported with MLP along with all per-VC WRED and LLQ and VC bundles on any of the aggregation platforms such as the Cisco 7200 or 7500?

A. Including MLP and PPPoA in the Parallel Express Forwarding (PXF) path is in development.

Q. What are the minimum VIP requirements to run fancy queuing on an OC-3 link?

A. You need at least a VIP2-50 to run fancy queuing on an OC-3 link.



Q. Is per-VC queuing available on ATM interfaces?

A. Yes. PA-A3s support per-VC queuing is available on ATM interfaces.

Q. **sh int** statistics show per-VC queuing enabled on PA-A3s port adapters, even when the fancy queuing is not configured. How can you revert back to FIFO?

A. Default per-vc queuing is **per-vc FIFO** and hence **show int** shows per-vc queuing. So if you don't configure service-policy or WRED, the VC queuing will be FIFO.

Q. Is there any difference in per-VC CBWFQ functionality for PA-A3-E3 cards?

A. No. PA-A3-E3 is exactly the same as the other types of PA-A3; only the framer chip is different. There shouldn't be any difference for per-VC CBWFQ functionality.

## Policing

Q. What is the difference between the Class-Based Policer and CAR? Which method is recommended for rate limiting?

A. The differences between the Class-Based Policer and CAR are given in Table 4.

**Table 4** Class-Based Policer vs. CAR Rate Limiting

Class-Based Policer	CAR
• Enabled within a policy using Modular QoS CLI	• Enabled explicitly on an interface
• Classification is mandatory	• Classification is not mandatory. Can do per interface rate limiting for aggregated CIR.
• Three action items for conforming/non-conforming traffic: conform, exceed, and violet.	• Two action items for conforming/non-conforming traffic: conform.
• Different options for each action for conforming/non-conforming traffic	• Different options for each action for conforming/non-conforming traffic
• Uses separate token buckets for Bc and Be	• Uses one token bucket for Bc and Be
• Configured using MQC within a class using police command	• Configured per interface using rate-limit command
• Recommended to do rate limiting	• Legacy rate limit feature

Q. What matching criteria do the Class-Based Policer and CAR support for classification?

A. Table 5 gives matching criteria of the Class-Based Policer and CAR.

**Table 5** Class-Based Policer vs. CAR Classification Support

Class-Based Policer	CAR
• Matching criteria stated in MQC format	• Incoming/outgoing interface
• Incoming/outgoing interface	• All/any IP traffic
• All/any IP traffic	• DSCP/IP precedence
• DSCP or IP precedence value	• Defined by rate-limit access list
• Standard or Extended source/destination access list	• MAC address
• IP RTP ports	• Defined by rate-limit access list
• 0-99 qos-group Ids	• Standard or Extended IP access list
• CoS value	• 0-99 Qos-group IDs
• Predefined class-maps	

## Class-Based Policer

## CAR

- MPLS experimental value
- Source/Destination MAC address
- NBAR protocols

Q. What are some of the action policies for conforming/non-conforming traffic offered in both rate-limiting mechanisms?

A. The action policies offered are given in Table 6.

**Table 6** Class-Based Policer vs. CAR Conforming and Non-conforming Traffic

## Class-Based Policer

## CAR

	• continue
	• drop
• drop	• set-dscp-continue
• set-clp-transmit	• set-dscp-transmit
• set-dscp-transmit	• set-mpls-exp-continue
• set-mpls-exp-transmit	• set-mpls-exp-transmit
• set-prec-transmit	• set-prec-continue
• set-qos-transmit	• set-prec-transmit
• transmit	• set-qos-continue
	• set-qos-transmit
	• transmit

Q. Are the Class-Based Policer and CAR DSCP compliant?

A. Yes.

Q. When does rate limiting occur, before or after fragmentation? Before or after queuing?

A. Rate limiting of traffic in the Class-Based Policer, CAR, or Shaping occurs before fragmentation or queuing takes place.

Q. When CAR calculates bandwidth usage in determining rate limit conformance, does it take into account L2 header and/or IP header length?

A. Yes. It does.

Q. If packets are remarked using policer within a policy, would the rate-limiting, queuing, WRED, or shaping configured under the same policy account for a new priority value, or will it still act based upon the original priority value? Would this differ any if the **set** option is used to remark the packets?

A. In all cases, it will act upon the original priority value. Common classification happens for all the features before the packets are marked by the policer. This means that any attributes changed by the policer (for example, the **setting ip precedence** command) is not considered for classification purposes. Policies defined on the next hop will be able to take advantage of the remarked priority. If you need to take actions based on the new priority, then mark the packets on the incoming interface. Then apply the policies on the outgoing interface based on this new priority.



Q. It is explained that the policer does not use buffers. What is the relationship between queuing and policing if the policer is configured with CBWFQ?

A. It is true that buffering comes with shaping and not policing. In the case that policing is configured with CBWFQ or other queuing mechanism, in the time of congestion, packets that conformed may need to be queued. This is when the defined queuing mechanism gets activated.

Q. Is the Class-Based Policer supported on a tunnel interface?

A. The Class-Based Policer can be applied to an interface or subinterface. It is not supported on Fast EtherChannel<sup>®</sup> technology, PRI, ATM LAN Emulation (LANE), or tunnel interfaces.

Q. Will packets destined to the router be rate limited with the inbound policer/CAR? Does traffic have to be CEF switched to work this way?

A. When packets are destined to the router, and if the inbound policer/CAR is enabled, packets will be rate limited. CEF is not required in this scenario. Remember that on a Cisco 7500 Series Router, traffic policing can monitor CEF switching paths only.

Q. Can locally generated packets be rate limited by the outbound policer/CAR? Does traffic have to be CEF switched to work this way?

A. Yes. No.

Q. If CIR is set to 45 kbps, does this mean that 45 bps is guaranteed during congestion and noncongestion periods?

A. CAR or class-based policer assures maximum rate limit and not minimum bandwidth to a class. To guarantee minimum bandwidth, use queuing (bandwidth or priority) options.

Q. If CIR is set to 45 kbps, does this limit apply only during congestion times?

A. No. Rate limiting will get activated as soon traffic exceeds CIR, regardless of congestion.

Q. What are default Bc and Be values used if I do not configure specific values? What effect does it have on token bucket?

A. Bc and Be are specified in bytes. Both will default to 1500 bytes. Make sure the maximum transmission unit (MTU) size configured on an interface is equal to or less than the default or configured Bc or Be values. Initial numbers of tokens in token buckets are equal to the size of Bc and Be. For example, if Bc = Be = 1800, then initially both token buckets will start full at 1800 bytes.

Q. Explain one-rate token bucket algorithm? How does the algorithm decide if a packet conforms or exceeds CIR? What happens when a packet arrives at the interface when the Class-Based Policer is applied?

A. One-rate token bucket accounts for only exceed action. When violate action is not configured, the one-rate token bucket algorithm is used. The following example explains how one-rate token bucket works:

```
policy-map POLICE
class onebucket
police 8000 1000 conform-action transmit exceed-action drop

interface fastethernet 0/0
service-policy output POLICE
```

The treatment of a series of packets leaving Fast Ethernet interface 0/0 is dependant on the size of the packet and the number of bytes remaining in the conform bucket.

The time interval at which the token bucket is replenished with tokens is  $T_c = B_c / CIR$ . Thus  $T_c$  in this example would be  $T_c = 1000 \text{ bytes} / 8000 \text{ bps} = 1 \text{ sec}$ . This result guarantees that 1000 bytes are added every second. Tokens are removed per number of bytes and not per number of packets that get transmitted!

In this example, the initial token bucket starts full at 1000 bytes. If a 450-byte packet arrives, the packet conforms because enough bytes are available in the token bucket. The conform action (transmit) is taken by the packet, and 450 bytes are removed from the token bucket (leaving 550 bytes).

If the next packet arrives 0.25 seconds later, 250 bytes are added to the token bucket  $(0.25 * 8000) / 8$ , leaving 700 bytes in the token bucket. If the next packet is 800 bytes, the packet exceeds and the exceed action (drop) is taken. No bytes are taken from the token bucket.

If no packet arrives for 2 seconds, only 1000 bytes for the first second are added to the token bucket, and an additional 1000 from the second minute are discarded.

Q. What is the two-rate token bucket algorithm? How does the algorithm decide if a packet conforms or exceeds CIR? What happens when a packet arrives at the interface when the Class-Based Policer is applied?

A. When the **violate-action** option is specified while configuring a policy with the **police** command in Cisco IOS Release 12.1(5) T onward, the token bucket algorithm uses two token buckets. Or in other words, we have conform and exceed token buckets. The following example uses the token bucket algorithm with two token buckets.

```
policy-map POLICE
class twobucket
police 8000 1000 1000 conform-action transmit exceed-action set-dscp-transmit 4 violate-action drop

interface fastethernet 0/0
service-policy output POLICE
```

The treatment of a series of packets leaving Fast Ethernet interface 0/0 is dependant on the size of the packet and the number of bytes remaining in the conform and exceed token buckets.

$B_c$  size tokens are replenished in the conform token bucket every  $T_c$ . Tokens are replenished in the exceed token bucket with the number of overflowed tokens from the conform bucket. If the silent period lasts long enough (to cover full intervals) and no more tokens are removed from the conform bucket,  $B_e$  size tokens are replenished in the exceed bucket at  $T_e$ . Tokens are removed per number of bytes and not per number of packets that get transmitted!

In this example, the initial token buckets start full at 1000 bytes. If a 450-byte packet arrives, the packet conforms because enough bytes are available in the conform token bucket. The conform action (transmit) is taken by the packet, and 450 bytes are removed from the conform token bucket (leaving 550 bytes).

If the next packet arrives 0.25 seconds later, 250 bytes are added to the conform token bucket  $(0.25 * 8000) / 8$ , leaving 700 bytes in the conform token bucket. If the next packet is 800 bytes, the packet does not conform because only 700 bytes are available in the conform token bucket. The exceed token bucket, which starts full at 1000 bytes (as specified by the  $B_e$  size) is then checked for available bytes. Because enough bytes are available in the exceed token bucket, the exceed action (set the QoS transmit value of 1) is taken and 800 bytes are taken from the exceed bucket (leaving 200 bytes in the exceed token bucket).



If the next packet arrives 0.40 seconds later, 400 bytes are added to the token buckets  $(0.40 * 8000)/8$ . Therefore, the conform token bucket now has 1000 bytes (the maximum number of tokens available in the conform bucket), and 100 bytes overflow the conform token bucket (because it took only 300 bytes to fill the conform token bucket to capacity). These overflow bytes are placed in the exceed token bucket, giving the exceed token bucket 300 bytes.

If the arriving packet is 1100 bytes, the packet does not conform because only 1000 bytes are available in the conform bucket. Similarly, the packet does not exceed because only 200 bytes are available in the exceed bucket. Therefore, the packet violates and the violate action (drop) is taken.

Q. Does CAR use the same algorithm described in the previous questions?

A. Although the basic concept of Tc and sending packets only if the tokens are available are the same, the algorithm for calculating and replenishing tokens is very different from that of Class-Based Policing as described in the previous two examples. CAR uses the concept of only one token bucket in both scenarios. Please refer to Q/A later in this section to learn more about how token algorithm for CAR works.

Q. What Bc or Be values are recommended? Should I know Tc or Bc or CIR? Where should I start?

A. CIR should be known for a class, and this information needs to come from the business requirement. Depending on traffic type, an approximate value of Tc should be defined as well. Keep Tc small if you are supporting delay-sensitive traffic such as voice. Based on these factors, calculate Bc. In other words, if the nature of the traffic is very bursty, expect long silent periods spaced by huge bursts, and configure Be high to deposit a lot of tokens in the exceed token bucket. Remember to change the MTU size on the interface accordingly; it should be greater than or equal to the burst size.

Q. Is hierarchical policing supported?

A. Yes.

Q. How can I verify that policing is working after it is applied to an interface?

A. Use the **show policy-map int <> EXEC** command to verify that rate limiting is active. Conformed, exceed, transmit, and drop counters keep track of all the matched packets:

```
Router# show policy-map interface
```

```
Ethernet1/7
```

```
service-policy output: x
```

```
class-map: a (match-all)
```

```
0 packets, 0 bytes
```

```
5 minute rate 0 bps
```

```
match: ip precedence 0
```

```
police:
```

```
1000000 bps, 10000 limit, 10000 extended limit
```

```
conformed 0 packets, 0 bytes; action: transmit
```

```
exceeded 0 packets, 0 bytes; action: drop
```

```
conformed 0 bps, exceed 0 bps, violate 0 bps
```

Q. Is per-VC input CAR supported on MPLS PE?

A. Per-VC input CAR with MPLS PE is indirectly supported by configuring input CAR and a VC under a subinterface (note that input CAR cannot be configured directly under a VC).

Q. On the Cisco 7500, is CAR supported on non-VIP interfaces?

A. CAR is supported on non-VIP interfaces; distributed CAR (DCAR) is not.

Q. Is CAR supported on BVI interfaces?

A. No.

Q. How is the token rate calculated?

A. In the context of CAR, token rate is calculated over the hardware clock. For the Cisco 7500 platform, it would be per 4 ms.

Q. With a following CAR statement, will the packets be dropped after the throughput exceeds 45 Mbps? How do the normal burst and Be values affect the drop rate?

interface serial4/0

rate-limit input 45000000 16000 24000 conform-action transmit exceed-action drop

A. Yes. In case of congestion, packets will start dropping after the traffic exceeds 45 Mbps. The behavior implied by normal and extended burst values is described in the token bucket algorithm:

As each packet has the CAR limit applied, tokens are removed from the bucket in accordance with the byte size of the packet. And tokens are replenished at regular intervals, in accordance with the configured committed rate. The maximum number of tokens that can ever be in the bucket is determined by the normal burst size. (So far this is just standard token bucket). Now, if a packet arrives and available tokens are less than the byte size of the packet, then the extended burst comes into play. If there is no extended-burst capability (extended burst can be achieved by setting the extended burst value to equal the normal burst value), then operation is as in a standard token bucket (that is, the packet will be dropped if tokens are unavailable). However, if the Be capability is configured (that is, extended burst > normal burst), then the stream is allowed to borrow more tokens (if under standard token bucket, there would be none available). The motivation is to not enter into a tail-drop scenario, but rather gradually drop packets in a more red-like fashion.

Q. Is there anything special you have to do to enable Turbo CAR, or is it used instead of regular CAR automatically when you configure a rate limit in 12.0(7) S or later code?

A. When the access control list (ACL) is compiled, it invokes the Turbo ACL code. Add **access-list compiled** to the configuration to invoke the Turbo ACL code. This feature is available for gigabit switch routers (GSRs) with Engines 0 and 1.

Q. When does CAR metering get invoked?

A. As soon as CAR is configured. Packets may or may not be dropped, depending on the CIR configured for each rate limiter.

Q. Does CAR give the best performance for coloring, or does policy-based routing work better?

A. In general, CAR has more work to do (that is, it must meter every packet as well, even if your need is to just mark them). PBR on the contrary, when set up with a **route-map** command for marking, will only mark the packets. Both are CEF supported, so PBR is at least as fast (if not faster) than CAR.

Q. Do CAR or the Class-Based Policer support aggregate and micro flow rate limiting?

A. Both CAR and the Class-Based Policer support aggregate rate limiting. The policer on the Catalyst® 6000 Switch supports both methods.



## Shaping

Q. What are the matching criteria for classifications supported for Class-Based Shaping and generic traffic shaping (GTS)?

A. Matching criteria for classification are given in Table 7:

**Table 7** Matching Criteria for Classification Supported for Class-Based Shaping and GTS

Class-Based Shaping	GTS
• Matching criteria stated in MQC format	• Outgoing interface
• Outgoing interface	• All/any IP traffic
• All/any IP traffic	• DSCP/IP precedence
• DSCP or IP precedence value	• Defined by rate-limit access list
• Standard or Extended source/destination ACLs	• MAC address
• IP RTP ports	• Standard or Extended IP access list
• 0-99 qos-group Ids	• Source/Destination MAC address
• CoS value	
• Predefined class-maps	
• MPLS experimental value	
• Source/Destination MAC address	
• NBAR protocols	

Q. Are both CB-Shaper and GTS DSCP compliant?

A. Yes.

Q. What are the recommended Bc and Be values?

A. Unlike that in CAR, it is recommended not to configure sustain burst value and let the algorithm calculate it. If excess burst is not configured, 0 is used by default (noticed from the router Bc = Be???). Both are configured in bits per interval.

Q. Where is it appropriate to use shaping?

A. Shaping is used to define the maximum rate that traffic within a class can be transmitted regardless of the presence of congestion on an interface. It is appropriate to use in the environment where:

- There is speed mismatch between central and remote sites
- If remote-to-central site oversubscription exists
- You need to prohibit bursting above committed rate
- You need to avoid instantaneous congestion and retransmission of TCP traffic and enforce the rate limit gracefully

Q. Is GTS supported on ATM interfaces? If yes, since what Cisco IOS version and on all ATM types? Is it recommended to use GTS on ATM interfaces?

A. On Cisco IOS Software-based platforms, GTS is supported on ATM interfaces from 11.2 onward. GTS is not supported on those ATM interfaces where per-VC shaping is done inside SAR. On numerous port adapters (PAs), such as Deluxe, VC shaping is done inside the SAR chip. On other PAs, it is done by default inside the software-based SAR. In these scenarios, placing another shaper on top of it doesn't make sense.

Q. What are some of the differences between GTS/CB-Shaper and FRTS?

A. Differences between GTS/CB-Shaper and FRTS are given in Table 8:

**Table 8** GTS/CB-Shaper vs. FRTS

GTS/CB-Shaper	FRTS
• Shaper for HDLC, FR, and ATM VCs.	• Shaper for Frame Relay Interfaces only
• Not supported per DLCI	• Supported per DLCI
• Support for class based, interface based, and group based	• MQC support on the roadmap.
• Supported via MQC	• Supports FRF.12
• Doesn't support FRF.12	

Q. What factors need to be accounted for if I want to do GTS and not DTS on the Cisco 7500?

A. To do GTS, packets need to go to the Route Processor (RP). Make sure that DCEF is not enabled. If DCEF is enabled, the packets will not go to the RP.

Q. In the case of shaping, if I am also remarking packets using the **set** option, do packets get marked first, and then get shaped (during congestion), or do they get marked before they are dequeued?

A. Packets are marked before they are shaped.

Q. Is FRTS necessary for per-VC queuing?

A. You need FRTS for per-VC queuing. For instance, if you want to implement per-VC LLQ because you have voice traffic, you need to enable FRTS.

Q. Is CEF required for the shaper to operate on locally generated traffic?

A. No.

Q. How can you verify that shaping is activated on an interface?

A. Router# show policy-map GTS\_in\_ModCLI

Policy Map GTS\_in\_ModCLI

Class shaped

Weighted Fair Queuing

Bandwidth 241 (kbps) Max Threshold 64 (packets)

Traffic Shaping

Average Rate Traffic Shaping

CIR 241000 (bps) Max. Buffers Limit 1000 (Packets)

Policy Map CBWFQ\_in\_GTS

Class cust\_A

Weighted Fair Queuing

Bandwidth 25 (%) Max Threshold 64 (packets)

Class cust\_B

Weighted Fair Queuing

Bandwidth 25 (%) Max Threshold 64 (packets)

Class cust\_C

Weighted Fair Queuing

Bandwidth 25 (%) Max Threshold 64 (packets)



Class class-default  
Weighted Fair Queuing  
Flow based Fair Queuing

WRED

Q. Why use WRED?

A. When congestion occurs, excess traffic is tail dropped. Tail drop treats all traffic equally and does not differentiate between classes of service. It causes global synchronization. Global synchronization manifests when multiple TCP hosts reduce their transmission rates in response to packet dropping, and then increase their transmission rates once again when the congestion is reduced. This scenario results in inefficient use of bandwidth. WRED is a congestion-avoidance technique that monitors the link and starts dropping fewer packets from selected traffic, avoiding bottlenecks and global synchronization on an interface. It is intended for TCP traffic. WRED calculates drop probabilities for a packet based on its IP Precedence or DSCP values and their associated minimum/maximum thresholds and mark probability values. This setup allows differential treatment for different traffic types.

Q. What is the function of minimum/maximum thresholds and drop-probability values?

A. The decision of when to start dropping packets when the queue reaches a certain height (**avg\_q\_size**) and when to stop dropping packets when the queue drains below certain value (**avg\_q\_size**) is based on the defined minimum/maximum threshold values. The drop-probability value associated with the packet priority value participates in determining the number of packets that should be dropped first out of the total packets that may be present in a queue.

Q. What are the default minimum/maximum threshold values, and what is the default drop probability?

A. The default minimum/maximum thresholds and drop-probability values for each IP Precedence and DSCP value is given in Table 9:

**Table 9** Thresholds for Precedence and DSCP Values

DSCP/IP Precedence	Min Threshold	Max Threshold	Drop Probability
0	20	40 for all the DSCP	10 for all the DSCP
1	22		
2	24		
3	26		
4	28		
5	30		
6	32		
7	34		
8	22		
9	22		
10	24		
: inc by 1'til...	:increment by 2...		
15	34		
16	24		
17	22		
: increment by 1...	: inc by 2..		
23	34.		

DSCP/IP Precedence	Min Threshold	Max Threshold	Drop Probability
24	26		
25	22		
: inc by 1	: increment by 2...		
31	34		
32	28		
33	22		
: inc by 1	: inc by 2		
39	34		
40	30		
41	22		
:inc by 1	:inc by 2		
47	34		
48	32		
49	22		
: inc by 1	: inc by 2		
55	34		
56	34		
57	22		
: inc by 1	: inc by 2		
63	34		

Q. What is the relationship between hold queue and threshold size? How can you calculate default minimum threshold values for various precedence or DSCP values?

A. Notice that the maximum threshold is equal to the default hold queue size 40 on an interface. Hold queue size is equivalent to the number of packets that can be held within a queue. The hold queue length range is 0 to 4096; the minimum/maximum threshold range is 1 to 4096. The default maximum threshold will reflect the defined hold queue size. In other words, if the hold queue is changed to 100, the maximum threshold will change to 100.

default max threshold = hold\_q\_size

default minimum threshold for precedence 0 = integer value of (max threshold/2)

If the precedence or DSCP value is increment by 1, then minimum threshold should be incremented by an integer value derived from (Maximum threshold/2)/9. For example, if the hold queue(out) is set to 100, the maximum threshold for precedence (or DSCP) 0 in this case is 100, so the minimum threshold is 50. With the same maximum value, for Precedence 1, the minimum threshold will be:

$(100/2)/9 = 5$ .

\*WRED treats non-IP traffic as traffic with Precedence 0 priority.

Q. What needs to be considered if you don't accept the default threshold values?

A. You can configure desired threshold values manually using **random-detect dscp** or **random-detect dscp** under a policy. Manually configured values will dictate as long as hold queue size is greater than the maximum threshold.

Q. How can I find out the hold queue size of an interface?

A. Use the **show run** or **sh queue <int#>** command. For example,

```
interface Serial4/0
mtu 1800
```



```
ip address 4.4.4.1 255.255.255.0
load-interval 60
service-policy output PSTGRE
hold-queue 20 out ---- other than a default value is displayed!
```

or

```
7200-UUT#sh queue ser4/0
Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 346628
Queuing strategy: weighted fair
Output queue: 0/20/64/346628 (size/max total/threshold/drops)
Conversations 0/1/256 (active/max active/max total)
Reserved Conversations 1/1 (allocated/max allocated)
Available Bandwidth 636 kilobits/sec
```

Q. What should I do if, even after changing the hold queue size on an interface, **sh polic-map int <>** shows old default threshold values?

A. Detach and reapply the policy on the interface.

Q. What is n, and what is the recommended value?

A. The factor n is used to calculate weight for mean queue depth. Default values are the recommended values, especially for exponential weight factor (the default is 9). Do not change this unless it is necessary. This command is available with different configuration options. For example:

In MQC:

```
7200-UUT(config-pmap-c)#random-detect exponential-weighting-constant <1-16>
```

Q. What are different ways of configuring WRED?

A. A WRED can be configured in three ways:

1. Class-Based WRED can be configured via MQC per class per policy. WRED will follow minimum/maximum threshold and drop-probability values configured for each DSCP/Precedence within a class. The policy containing the WRED configuration can be applied to one or multiple interfaces.
2. Interface-based WRED can be configured directly on an interface.
3. Group-based WRED can be applied to multiple interfaces. It includes a list of DSCP/Precedence precedence values and associated minimum/maximum threshold and drop-probability values.

Q. Which classes can WRED be applied in?

A. WRED can be applied in all the Assured Forwarding (AF) (RFC 2597) classes defined in the Diffserv standard; in other words, all the user-defined classes as well as the best-effort (class-default) class.

Q. Can WRED be applied in conjunction with Class-Based Policer, CBWFQ, and Class-Based Shaper?

A. Yes. Remember that WRED works on software queues. It will work on CBWFQ or WFQ. Make sure a class has queuing enabled in conjunction with the policer or shaper.

Q. How quickly does WRED respond after traffic goes above the minimum threshold to start dropping packets, and how quickly does it slow down when queue drains to or goes the minimum queue threshold?

A. WRED may be slow to start dropping, and it might even continue to drop after the queue size has fallen below the minimum threshold. Lower and upper bound limits are applied to the calculated average queue size, which depends on exponential weighting constant and the current and old average queue size.

$$\text{avg\_q\_size} = (\text{old\_avg} * (1 - 1/2^n)) + (\text{current\_queue\_size} * 1/2^n)$$

If you look at the equation, *old avg queue size* is the leading factor in determining average queue size even if  $n = 1$  (mix value) and the current queue size is 256 (for a 2M link). This helps avoid drastic changes in average queue size, a setup that translates into efficient bandwidth utilization. If you start dropping packets too soon, you will underutilize the bandwidth.

The *old\_avg* is the average computed the last time a packet was enqueued. It starts out with 0 initially. The *current\_queue\_size* is the instantaneous queue depth.

Initially:  $\text{old\_avg} = 0$ ,  $\text{current\_queue\_size} = 0$

pkt arrives:  $\text{avg\_q\_size}(\text{computed}) = 0$ , packet gets enqueued,  $\text{current\_queue\_size} = 1$  next packet arrives,  $\text{old\_avg} = \text{avg\_q\_size} = 0$  so  $\text{avg\_q\_size} = 1/2^n$

lets say packet gets enqueued:  $\text{current\_queue\_size} = 2$

next packet comes in:  $\text{old\_avg} = \text{avg\_q\_size} = 1/2^n$  thus  $\text{avg\_q\_size} = 1/2^n * (1 - 1/2^n) + 2 * 1/2^n$

Q. How does WRED work? What happens when a packet arrives at the interface?

A. When a packet arrives and WRED is enabled, the following events occur:

1. The average queue size is calculated using  $\text{avg\_q\_size} = (\text{old\_avg} * (1 - 1/2^n)) + (\text{current\_queue\_size} * 1/2^n)$  (Where  $n$  could be from 1 to 16; default is 9).
2. If the average is less than the minimum queue threshold, the arriving packet is queued.
3. If the average is between the minimum queue threshold and the maximum queue threshold, the packet is either dropped or queued, depending on the packet drop probability.
4. If the average queue size is greater than the maximum queue threshold, the packet is automatically dropped.

Q. Is WRED on by default on an interface?

A. No. Tail drop is on by default.

Q. Does WRED act on the remarked priority or the original priority of a packet?

A. WRED acts on the remarked priority value, regardless of whether packets are remarked within a class using the **set** option or the policer.

Q. Does WRED act on software queues as well as on the hardware queue?

A. WRED always acts on the software queue, and never on the hardware (interface) queue.

Q. What Cisco IOS version does WRED honor DSCP in?

A. 12.1.5T and above.

Q. Does WRED act on EXP bits in an MPLS network?

A. Yes. This is supported.

Q. Is CEF required for WRED to operate on locally generated packets?

A. No.

Q. Why can't I configure WRED on a one-port Gigabit Ethernet card in GSR?

A. WRED is not supported on this card.

Q. Is WRED supported on all port adapters or on only a limited number of them?

A. Per-VC WRED is supported on select port adapters only, such as the PA-A3 for the Cisco 72xx and 75xx. Per-interface WRED is supported on the PA-A1 as well; it is not supported per subinterface.



Q. Does WRED drop packets from fancy queues?

A. Prior to 12.1, WRED did not work on PQ, CQ, or WFQ. Starting from 12.1, it works with CBWFQ.

Q. Please explain the differences between Modified Deficit Round Robin (MDRR) and WRED.

A. (M)DRR is for emission priority or bandwidth allocation. It works only when queues build up, in other words, during congestion. DRR or MDRR is needed only if you want to allocate more bandwidth to some higher-priority traffic (for example, voice) or latency-sensitive traffic. If all the traffic is of the same priority to you, no MDRR is needed. It is not a congestion-avoidance tool, which WRED is the target for. (M)DRR works best together with WRED. If MDRR is applied on to the fabric, it can possibly affect packets destined to different output interfaces, but if it is applied on the fabric, it will affect only the packet on that specific interface.

Q. Does MDRR imply implementation of the high-priority queue? Likewise, does DRR imply no high-priority queue?

A. Algorithmically, DRR describes a scheduling algorithm among a set of queues or queuing systems. In a hierarchical queuing scheme, it doesn't preclude one somewhere else in the hierarchy. For example, you could use DRR within a class of, say, seven queues, and have a hierarchical scheduler that prioritized an eighth queue—or several such queues, perhaps with a stochastic WRR or DRR per-flow scheduler inside each—over that DRR system.

#### VoIP-Related Issues

Q. What are your general QoS recommendations for VoIP traffic?

A. Generally:

- Use 12.0.7T or later to get the best queuing features.
- Do not use VoIP on a PVC that also carries voice over Frame Relay (VoFR).
- Set the IP Precedence to 5 on the dial peer.
- *Do not* use WRED on voice queues.
- *Do not* mark voice packets as DE.
- Turn on dual tone multifrequency (DTMF) relay for low-bit-rate coders/decoders (codecs) (8K and below).
- Set echo, loss/gain parameters according to network loss plan.
- If TCP delays affect DTMF-relay performance, use **cisco-rtp** for DTMF-relay.
- Measure network packet delay—the goal should be 150 to 200 ms one-way delay

#### Queuing:

- LLQ—Classify voice in a “priority” class.
- Set the bandwidth of the voice class to the aggregate voice bandwidth on the link or VC (and allow for a little overhead).
- If LLQ is not available, use “IP RTP Priority” (LLQ is preferable to IP RTP priority and is reserved).

#### Bandwidth:

- Set bandwidth using the **priority** statement in the LLQ configuration (or in the IP RTP Priority statement).

#### Fragmentation (for link speeds < 1.5M):

- Fragment to 10-ms delay—optimize size for backbone characteristics.
- Set fragment size so that voice packets do not get fragmented.
- For leased lines, set **ppp multilink fragment-delay** on the interface.
- For Frame Relay:
  - Set **frame-relay fragment** in the Frame Relay map class.
  - Fragment all PVCs carrying data on the interface if at least one PVC carries voice.

**Traffic shaping (if Frame Relay is used as a Layer 2 technology):**

- Configure FRTS on the interface.
- Set Bc to 10 ms (1/100) of CIR.
- Set **mincir** >= to voice bandwidth (if adaptive shaping is used).
- Shape strictly to CIR one PVC carrying voice; don't configure burst.
- Shape both sides of the VC to prevent egress blocking.

Q. With voice over IP Security (IPSec), how much latency do Data Encryption Standard (DES) and Triple DES (3DES) encryptions add? What is the quality?

A. The IPSec overhead is not that bad, assuming you have enough bandwidth. As far as the latency is concerned, after the VPN is actually set up, it is about the same as without encryption.

If you are running T1s and T3s, you should be fine. However, there is more latency when the virtual private network (VPN) is building. After it is set up, the latency is the same as on any other VoIP call.

Q. How would you calculate a serialization delay for a voice packet stuck behind a data packet if you cannot fragment data packets?

A. For a data packet of 1500 MTU and 256-kbps egress speed, Delay =  $(1500 \times 8) / 256000 = 46.9$  msec.

Q. Does link fragmentation and interleaving (LFI) apply to ATM to provide guaranteed delay for VoIP traffic?

A. No. LFI applies only to Frame Relay. By specification, when a packet enters the SAR in ATM, no other packet can preempt. SAR is not good enough for LFI; it does a form of fragmentation, but it does not do interleaving. Interleaving is what voice packets need to avoid serialization (or VC traffic shaping "serialization" in this case) delays.

Q. What other protocol can be used to support fragmentation/interleaving over ATM and Frame Relay?

A. You can use PPP/MLP over ATM and then use LFI for PPP. This will be available in 12.1.4/5T. MLP is already supported over Frame Relay and ATM. LFI support for MLP over ATM and MLP over Frame Relay and interworking between MLP over ATM and MLP over Frame Relay are in development.

Q. How can you overcome low-speed ATM PVCs (768kbps or lower speed) limitations and guarantee serialization delay for VoIP traffic?

A. Use a bundle PVC feature that facilitates VoIP constant bandwidth requirements by dedicating some VCs for voice and others for data traffic.

Q. Is the CompressReal-Time Transport Protocol (cRTP) supported over ATM?

A. No.

Q. Does VoFR on the Cisco 2600/3600 support multiple subframes within a frame? Is the use of one PVC for both voice and data on the Cisco 2600/3600 allowed? (One of the 255 subchannels can be used as data protocol in FRF.11). Is it necessary to configure a specific subchannel for data?

A. Voice and data are supported on the same DLCI.

You do not need to worry about dedicating a specific subchannel; it will do this automatically. Cisco does support 255 context identifiers (CIDs) (FRF.11 subchannels) and you can, if you wish, explicitly configure the CID to be used for data. The default is 4.



However, Cisco does not support the ability within FRF.11 to put multiple “CID payloads” into a single Frame Relay frame, that is, to transmit a Frame Relay frame that has inside it a frame belonging to CID 4 and a frame belonging to CID 5, and so on). In the Cisco implementation, each CID frame is also a separate Frame Relay frame. But it doesn’t matter; it might provide a slight optimization on the number of frames or bandwidth used, but it is not an issue that matters to network design anyway.

Q. When a PVC is configured for VoFR only, do we need to configure data fragmentation on this PVC?

A. No. But you still need to configure it on the other PVCs that share the same interface with this one.

Q. When you configure FRF.12 on one PVC, does this PVC get used **only** for data? If so, are we going to configure this PVC for data with FRF.12 (and with one map class), and the other PVCs for voice (and with other map classes) in the same interface?

A. You need to have a separate map classes for the voice PVC and the data PVC.

You do not configure FRF.12 for a PVC. You configure fragmentation *only* for a PVC. The software will figure out whether or not it should use FRF.12 (data PVC) or FRF.11 Annex-C (VoFR PVC).

All PVCs that carry data and share an interface with at least one PVC carrying voice must be fragmented if the speed of the VC is less than 1.5M.

Q. Given a VC bundle of two VCs, one non-real-time variable bit rate (VBR-nrt) for VoIP and one unspecified bit rate (UBR) for data, will the VBR-nrt get its guaranteed bandwidth defined by the sustainable cell rate/peak cell rate (SCR/PCR) and UBR get everything else up to the PCR?

A. Yes.

Q. Are the router ATM interfaces that support VC bundles sophisticated enough to interleave packets between different VCs? Can VC bundles be used as an interleaving technique to prevent head-of-line blocking by big data packets on the UBR PVC?

A. Yes. You can have one VBR-nrt (PCR = SCR) for voice and one VBR-nrt (pcr! = scr) for data if you want data to get its guaranteed bandwidth as well. PA-A3 can even allow you to raise the transmission priority of voice VC when it is a low-speed bundle.

Q. Can **IP RTP priority** be applied per VC?

A. Yes. **IP RTP priority** can be enabled per VC along with FRF.12 (fragmentation).

Q. Are there any limitations about using the Cisco 2600/3600 with VoIP applications *and* VPN?

A. If you encrypt only data traffic and not voice traffic, there are no interactions, and you can do that without problems. As far as tunneling voice, there will be interactions between IPsec and voice. The bottom line is that there’s an FIFO queue for traffic entering the encryption engine, and this causes jitter in voice. Further, to ensure the voice QoS features on the WAN interface, make sure you use tunnel pre-classification feature to copy the packet priority value to the tunnel header. You will also lose the ability to do LFI on the WAN link in a sense that it will fragment, but it cannot interleave because everything is encrypted.

QoS for VPNs

Q. How can a value present in the ToS field (for IP packet) be preserved if a packet is classified by the IP Precedence of DSCP value after a packet is encrypted for a tunnel?

A. Use the tunnel preclassification feature. This will copy the ToS value of an original packet into the tunnel header to preserve the QoS associated with particular DSCP or IP Precedence values. The same command applies to generic routing encapsulation (GRE), Layer 2 Tunneling Protocol (L2TP), or IPsec tunnels.

For example:

#### GRE and IPIP Tunnels

```
spider(config)#interface tunnel0
spider(config-if)#qos pre-classify
```

#### L2F and L2TP Tunnels on LNS

```
spider(config)#interface virtual-template1
spider(config-if)#qos pre-classify
```

#### IPSec Tunnels

```
spider(config)#crypto map secured-partner-X
spider(config-crypto-map)#qos pre-classify
```

Q. On which interface should I apply my policy, tunnel, or a physical interface?

A. If you want to allocate QoS based on information contained in an original packet, apply a policy to a tunnel interface. If you want to allocate QoS based on information contained within a tunnel header, apply it to a physical interface.

#### Miscellaneous

Q. Is MLP CEF switched with and without fragmentation on Cisco 36xx, 72xx, and 75xx platforms?

A. No. The current CEF code does not install “adjacencies” to multilink bundles. IP switching to MLP bundles must be fast switched instead (in IP terms, that’s using the **ip route cache** command instead of CEF).

Q. How do I measure packet loss through the network with regard to different classes of service?

A. Use Class-Based Queuing over Security Management Information Base (CBQoS MIB (available in 12.1(2) E, 12.0(12) S, and 12.1(4) T), SAA/IPM or QPM.

Q. What queue does ISIS updates/hellos use on Eng 0 and Eng 2 cards using MDRR?

A. Depending on the queue you map precedence value of 6. It only matters on the Tx since the PDU packets comes from GRP.

Q. What MIBs are available for QoS?

A. The new Cisco Class-Based QoS MIB provides you the same stats as the **sh policy int** command provides. It is available for VIP-based QoS (that is, the Cisco 7500 only) on 12.0(12) S as well as 12.1(2)E. It will be available for all other Cisco IOS Software-based platforms in 12.1(4/5) T.

Q. Is DSCP support available for Border Gateway Protocol (BGP) policy propagation?

A. Not yet.

**Corporate Headquarters**

Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
[www.cisco.com](http://www.cisco.com)  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 526-4100

**European Headquarters**

Cisco Systems Europe  
11, Rue Camille Desmoulins  
92782 Issy Les Moulineaux  
Cedex 9  
France  
[www.cisco.com](http://www.cisco.com)  
Tel: 33 1 58 04 60 00  
Fax: 33 1 58 04 61 00

**Americas Headquarters**

Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
[www.cisco.com](http://www.cisco.com)  
Tel: 408 526-7660  
Fax: 408 527-0883

**Asia Pacific Headquarters**

Cisco Systems Australia, Pty., Ltd  
Level 17, 99 Walker Street  
North Sydney  
NSW 2059 Australia  
[www.cisco.com](http://www.cisco.com)  
Tel: +61 2 8448 7100  
Fax: +61 2 9957 4350

**Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the**

**Cisco.com Web site at [www.cisco.com/go/offices](http://www.cisco.com/go/offices).**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China • Colombia • Costa Rica • Croatia • Czech Republic • Denmark • Dubai, UAE  
Finland • France • Germany • Greece • Hong Kong • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The  
Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia  
Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

All contents are Copyright © 1992—2001 Cisco Systems, Inc. All rights reserved. Cisco, Cisco IOS, Cisco Systems, and the Cisco Systems logo are registered trademarks of Cisco Systems, Inc. or its affiliates in the U.S. and certain other countries. All other brands, names, or trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0102R)

3/01 LW2057